06152025 Microsoft AI Tour Amsterdam Satya Nadella

Microsoft AI Tour
Satya Nadella, CEO & Chairman, Microsoft
Callie August, Director, Product Marketing, Microsoft

Joris Schoonis, CEO Microsoft Netherlands
Ronlad den Elzen, Chief Digital & Technology Officer, The Heineken Company
Mark Chaban, VP Commercial Cloud Solutions, Microsoft EMEA

Theater Amsterdam - Danzigerkade 5
1013 AP Amsterdam North Holland Province, Netherlands
Monday, June 16, 2025

SATYA NADELLA: All right, good morning.

(Applause.)

Good morning. It's fantastic to be back here in Amsterdam, especially in times like this, when we're all busy with a new platform shift. It's great to have a chance to share what's happening, get to see what's happening right here in the Netherlands. It's an exciting time, for sure.

One of the things, whenever you think about these platform shifts in tech, it's best to start with, what is that source of change? What's really driving, in some sense, these rapid changes we are seeing in AI capability? It's always, never one thing. This has always been true, whether it was the web or even back in the day with client server, or what was true in cloud. It's multiple S-curves all working together.

We have, in some sense, Moore's Law back again in great force, really helping each year double capability. On top of it, the system software improvements are really building on that, and then the model improvements are building on that. It's the combination of these three S-curves that are essentially leading to what is doubling of capability every three months or so. I mean, this is not something I've, myself, seen in the last 35 years that I've spent in tech. This is probably the most rapid increases in capability.

Now, how it manifests in even how we use AI, I mean, think about where we started, even couple of years ago, to where we are. We first started by saying, let's use AI to ask it questions, and it comes back with answers. Then we said, let's assign it tasks, and it completes tasks. And now, we are at the threshold of essentially having, for example, in the case of GitHub, a peer programmer, not just a pair programmer.

We are really getting to a place where we have all these three modalities of how we use AI, all simultaneously part of any product, whether it's Copilot or GitHub Copilot or your favorite AI tool. That's what we describe as the open agentic web. This idea that you can

have AI power all of these capabilities and build them out is the world that we are all collectively shaping.

Now, for us, it's not about technology. At the end of the day, it is all about taking that technology and ultimately, testing it and with our mission, which is, can we use it to empower every person and every organization on the planet to achieve more. To us, that's the goal, and of course, the business outcomes that it drives are also very important.

Any organization will really measure its adoption of AI by saying, is it enriching my customer experience? Is it enriching my employee experience? Am I able to change the business process, productivity for any business process internally? Are we able to even improve our innovation curve by using AI? These business outcomes ultimately are going to be the real test of productivity for AI.

Now, what I want to do for the rest of my time is just to walk through the tech stack that is being used by all of you, in fact, to build solutions that really take advantage of all of these advances in AI. And let's start at the very foundation, which is the infrastructure.

It's exciting to see the buildout of compute. I always think of, whether it's for a company or a country, tokens per Euro, per watt becomes the most important consideration. To that end, we're really making computing ubiquitous. We are building out Azure as the world's computer. We have 70+ data center regions, more than any other cloud provider around the world.

Right here in Europe, we have data centers in 16 countries. We are going to expand our compute capacity in Europe by 40% in the next two years. It's significant buildout of data center capability. Obviously, in Netherlands, we have massive Azure region, which we'll continue to expand as well. We're excited about that, but the thing that we're also very, very focused on is to make sure we're building out our cloud infrastructure in Europe, for Europe, and making sure that sovereignty is a top level commitment for us.

In fact, in April, we announced a set of European digital commitments. This spanned everything from the commitment we have to build a local AI ecosystem, a cloud ecosystem, a commitment to ensure that there is resilience, no matter what. Even any geopolitical instability or volatility, still there needs to be resilience, and we have our commitments to that. Protecting privacy, that's another commitment, cybersecurity, as well as making sure that we are strengthening Europe competitiveness with technology, including open source that's being built in the continent.

We have these commitments we made in April. And today, we want to build on it, when it comes to sovereignty. Microsoft has a very comprehensive view and approach to sovereignty. It spans sovereign public cloud. It also extends to sovereign private cloud as well as our national clouds that we are building.

Let's just dive right in and talk about what is some of the new announcements we are really pleased to be making today, when it comes to sovereign public cloud.

It starts with Data Guardian. Now, Data Guardian is about making sure that administrative operation on the cloud in Europe is done by employees of Microsoft, who are domiciled and resident in Europe. And all of the operations themselves are logged in an immutable log that is auditable. This is a guarantee on all of the administrative operations.

The second thing that we're also very pleased to announce is the external key management. That means you will be able to encrypt your data and have the keys completely resident in your private cloud. None of the key material is in the public cloud at all.

The combination of this Data Guardian and key management essentially gives you a fully sovereign cloud operation using, while at the same time, the public cloud. And we have made all of this possible with a fantastic set of controls in this regulated environment management, so any customer, any tenant, will have one dashboard to go to be able to see what's happening with the Data Guardian and manage that as well as the key management. We are really ensuring that you have, even in the public cloud, a fully sovereign operation that you can take advantage of.

But we're not stopping there. We're going to the next step and saying, okay, what's the way for us to support a fully sovereign private cloud. And this is where we are very excited to extend Azure Local with virtualization support. That means all of the containers and VMs that you have in the Azure public cloud are now available on Azure Local, but on top of it now, we will have Microsoft 365 Local. That means all of your productivity, collaboration software can run as well, completely in your private cloud.

We think that the combination of sovereign services for public cloud and sovereign private cloud really give you all of the technology to be able to then drive European cloud infrastructure with all of the sovereignty requirements. But on top of it, we also have local partners.

Right here, we have InSpark, for example, in Netherlands, who is going to be able to use our various services, whether on the public cloud or on the private cloud, to provide local services. And we have many other European partners as well who are going to help us make sure that we deliver on all of the sovereignty requirements by country in Europe.

We are very excited about all of these announcements. That's the first stanza of what we're doing, which is all around infrastructure.

The next piece, of course, is the AI platform. When it comes to building on top of all of this infrastructure, the first consideration is all around data, because there is really no AI without data. And to that end, what we are doing is we are building out the data infrastructure very comprehensively. You have all of the operational stores, whether it is SQL, Postgres or Cosmos. We have all of the analytical stores with Fabric. And then we also bring everything from Oracle or from Databricks or Snowflake, and data governance

with Purview, so a very comprehensive approach to how to rendezvous all your data in the cloud and make sure that you're ready for being able to take it to the next generation of AI compute that is just close to data. That's the first piece.

Then the next big thing that we're really investing in is the Azure AI Foundry. The best way to think about AI Foundry is this is your factory for building agents. If you think about any of the platform shifts we've had in the past, whether it was to the web or to the cloud, you always had an app server. We had an app server for building out cloud-native applications. We had an app server for building web-native applications. And now, you have an app server for building AI-native applications.

You have everything, starting from the models. We have 1,000+ models available in Foundry. We have everything from Mistral to OpenAI to Black Forest labs to Grok to DeepSeek. All of these models are available because going forward, pretty much any application you build will be multi-model. And one of the cool capabilities of Foundry is that you can provision the throughput once and use it with any model. It's not about throughput that you provision in Foundry is dedicated to one model. It spans all models. And that, I think, gives the flexibility to app developers, going forward.

Now, so you have Foundry with all of the app services, the models so that you can now build your agents, but to build it, you still need tooling. And this is where Microsoft always has been a tooling company. From 1975 to 2025, we've done one thing, which is start with developer tools, and we continue to build out our developer tools around it.

And what we are doing is with GitHub and VS Code and the Foundry SDK and Copilot Studio, you have all of the tools that you need in order to build out all these agentic applications. And in Netherlands, we have great momentum. With GitHub, we have 1.5 million GitHub developers, and increasing by 20% year over year. It's great to see the developer ecosystem really thrive and grow right in Netherlands.

And in fact, when we think about the tooling that we have built in to GitHub, this is the first product we built for the age of AI, back now three years ago, which seems pretty ancient. We started with code completions. We then added chat so that you can ask questions. Then we introduced the Agent Mode so that you can then give it assigned tasks.

And just last month, we launched our coding assistant, the GitHub Copilot Coding Agent. And this is the autonomous coding agent that just goes and picks up at issue and starts to work on it in its own branch. And so, I wanted to roll the video to give you a flavor for the coding agent.

(Video segment.)

SATYA NADELLA: That's pretty cool. It's the first time I've seen that video. It's pretty nice for this theater, I must say. It's got good music. (Laughter.)

(Applause.)

When you think about all of these platforms and tools, ultimately, it's the apps that people build. I had a chance – in fact, there's many, many names on the slide that you see right here, people who are building out these applications. And this morning, I had a chance to meet with some developers and get a sense for how people are building these applications, and it's exciting.

I met the team from ASML that's building out this multiagent toolkit, essentially, for giving everyone who works, 15,000 people inside of ASML who are trying to keep up with all of the procedures that they need to follow, which get updated, I believe, twice a year, with all the instructions. It's pretty precise stuff. They've built out this pretty sophisticated – in fact, they started with GPT-3.5 They told me it kind of didn't really work. But with now 4o, 4.1, it's working fantastically, and they've built this out. And it's just impressive to see these multiagent frameworks being pulled together to build out pretty sophisticated applications like this.

I had a chance to meet with both ChipSoft, which is like the EMR system that's got most of the hospitals in the Netherlands, along with UMC Utrecht. And one of the cool things, I mean, if you think about one of the greatest benefits of AI, it probably will be in healthcare, because a lot of the cost of healthcare is all in workflow.

Something as simple as discharge takes a lot of time. What they did is basically said, here's a prompt that goes to the data tier of the EMR system and does the discharge with one click. And in fact, the hospital system there that even open sourced the prompt so that all the other hospitals in the Netherlands can use it as well. That type of activity, once it kicks into gear in an ecosystem, can have really profound impact.

I had a chance to meet the City Archives team, which is fun. I guess you all will be celebrating your 75th birthday soon. They are taking what is 50 kilometers of archives and digitizing it, making it available with a beautiful frontend, with even a digital avatar that you can talk to and that will explain everything from the history of this great city.

Heineken, they have a vision of saying they want to be the best connected brewery. And they have their own software engineers using things like GitHub Copilot and completely modernizing their own software development lifecycle. They explain how they've taken some inspiration from even some of the work we have done in our own engineering systems and building it out right in Heineken. It's just what's fantastic to see.

And lastly, I had a chance to meet with the developers at Albert Heijn, who they've built out their application. I mean, they took their mobile app and introduced an agent called Steijn in it to be able to really help with their business. I just wanted to roll the video to give you a flavor for what they've done.

(Albert Heijn video segment.)

SATYA NADELLA: Yeah, it's great to see, again.

(Applause.)

It's great to see the maturity of the app stack manifest in the ambitions of the developers to be able to build out these applications. That's when you know that this is becoming a lot more mainstream. It's no more about one app, but it's about all the apps that we use suddenly having rich AI capabilities. And I think that that's when you really know a platform is really being born.

The last thing I want to talk about is, in fact, our tools and our applications for knowledge workers, which is Copilot. It's very exciting to see the launch of the new Copilot. The best way to conceptualize Copilot is it's five things in one.

Of course, it starts with chat. This is a place where you can have that simple request response, and it really allows you to get information, not only from the web, but from the work. It synthesizes. Even this morning, when I was doing some of the research on some of these things that I was going to talk about, I was able to just get the latest data from within Microsoft, which combined the internal data, which is all in M365 and web, and have it all synthesized by the LLM.

The second thing is you have search. You still want to be able to get all of the information, again, from the web, as well as inside, whether it's in Teams, whether it's in SharePoint, whether it's in email. You want to be able to search. You now have even search right there.

The next thing you have is the ability to create things, so you can prompt a PowerPoint slide deck, or you can, in fact, take something like a Word document and convert it into a PowerPoint slide deck. All of the creation tasks are now AI driven, so that module is built in to M365 Copilot.

The next thing is perhaps one of the more important concepts that are going to become as important as the SharePoint was back in the day, which is Notebooks. You want to be able to collect these heterogeneous collections of data. You want to be able to take, in fact, all of your chats, put them in. You want to take your documents, put them in your searches, whatever, and then use all that as grounding data for your project. Notebooks is an organizing layer for projects with heterogeneous data that are then used for grounding of your LLMs.

And then the last thing is agents. There are two agents that we have built, which we are very excited about. One is Researcher; the other is Analyst. This is about being able to have, effectively, someone who's like a research assistant that is there 24 by seven for you to be able to assign tasks, so that they can go off and do these asynchronous tasks. And Analyst is like having a data science person with you, 24 by seven. You can give it a couple of spreadsheets. It'll come back with insights, visualizations and so on.

This is the five things in one application. And now, of course, with Teams, think of it as the multiplayer version of this. You are using Microsoft 365 Copilot as a single player as your personal tool to be able to do your knowledge work with AI. And Teams is where you're working and collaborating with your colleagues as well as AI is right in place inside of your channel. Any agent you build, whether it's Researcher or Analyst, any custom agent, they'll all be available inside of Teams as well.

Copilot Studio is a place to build agents. In fact, our vision is very simple. Just like back in the day, you were able to build an Excel spreadsheet or a Word document, we want you to be able to use Copilot Studio to build agents. It's no more mystical than that. Here, there's an IT agent that just got built by grounding it with some SharePoint data, giving it some instructions, and then you have essentially created an agent. That's the idea of Copilot Studio.

And the impact of all of this is, at the end of the day, the productivity stats. At Microsoft, we are seeing things like customer support, where we are now 12% faster. Marketing, we're seeing improvements over 20%. In terms of conversion rates, our sales Copilot is really helping improve the yield of all of our sales teams by close to double digits. We have IT, when it comes to all of the IT services, massive, massive gains in terms of the deflection rates there with something like 35% or so increase in efficiency. And all of this is obviously being translated into productivity gains across our customers.

Now, we see great momentum. BDO in accounting right here, Philips, Rabobank, KPN, many, many others. In fact, I had a chance to meet – one of the cool things is I had a chance to meet one of our partners, Let's Copilot. And what they're doing is building out one of those Teams agents to help customers adopt Copilot and get the benefits of Copilot. Think of it as a gamified agent that helps people learn about Copilot, use Copilot on a daily basis across the entire Microsoft 365 system.

It's exciting to see all of this, so I wanted to show you the progress we are making with Copilot in its entirety. To do that, I want to invite up on stage, my colleague, Callie.

Callie, take it over.

(Applause.)

CALLIE AUGUST: Thanks, Satya. Our vision for the future of work is a Copilot for every employee and an agent for every business process. Let me show you what that means.

I start every morning in the Microsoft 365 Copilot app. This is your hub for navigating the new future of work with AI. Here, I have access to Copilot Chat across web and work data. I can search for information, and I can interact with agents like Analyst, Researcher and Create.

Now, I've been so inspired by seeing everyone bicycling around Amsterdam. I'm going to take on the persona of a marketer at a cycling company for these demos.

The first thing that I'm going to do is I'm actually going to ask Copilot Chat to outline my key emails from the last couple days and sort them in priority order. Copilot is going to comb through my email inbox and do that for me.

You can see here, it's outlining all of those emails, and it looks like our sales team has actually just put together a new analysis. I'm going to dive into that first.

To do that, I'm going to go over to our marketing and sales advisor agent. Now, I made this agent because a really key part of my job is identifying sales trends and turning them into marketing opportunities. Let me show you how easy this was to create in Copilot Studio.

I'm able to create agents in Copilot Studio using just my own words. I did that. Here, I gave it a little bit of a description about what I was looking for. I wanted it to identify key trends and provide marketing recommendations. Then I gave it some key instructions, like I wanted it to take seasonality and time of year into account in its recommendations. Lastly, I grounded it in the knowledge it would need in order to provide me accurate recommendations. You can see that here, and including that weekly sales analysis that's always updated live by the team.

Now, if I go back into the agent here, we can see it in action. I'm going to ask this agent to recommend a couple different bike models that I can use for this upcoming marketing campaign that I want to plan. It's going to look through that knowledge that I grounded it in, and it's going to take all of those instructions and the description of the agent into account.

And you can see here, it's given me two great examples. I've got a road bike and a mountain bike. Since we're just getting kicked off with summer, I'm going to lean into the mountain bike idea.

And I'm going to jump over to Researcher agent here, and I'm going to use our Researcher agent to actually help me create a full marketing plan for this campaign. I'm going to give it a little bit of direction, like I want it to look into competitive insights and our past campaigns. And Researcher is going to start thinking through the best way to tackle this problem.

Just like a seasoned colleague on my team, it's actually going to ask me some follow up questions. I'll give the Researcher agent, a little bit more direction, and it'll get started.

What Researcher is going to do now is an iterative process of looking through data on the web as well as my internal work files in order to pull together this marketing plan. It's going to go through a chain of thought reasoning process, which you can see it starting to do that right here.

And now, sometimes this chain of thought reasoning takes a couple of minutes. I'm actually going to jump ahead to the end result, but you can see that all populating live here, which is just so exciting to see.

We're going to go over to the result right here, and you can see that completed chain of thought reasoning at the top here. Then if I scroll down, you can see the full marketing plan that it built out. This is just incredibly thorough and so incredible to see live. And now, the best part is I can turn this into a page to collaborate with colleagues, or easily turn this into a Word document.

Now, the last part of this campaign is getting organized. For that, I want to move over to Copilot Notebooks. Now, this is a dedicated workplace for me to store all of my campaign files, all of the information that I need to stay organized on this campaign. And I'm able to interact with Copilot right here, based off of this data in my notebook. If I need some quick ideas, like a social media strategy or some social media ideas, I can do that right in the notebook, and Copilot will use those specific files to generate its answer.

And you can see here, it's giving me a little bit of strategy. It's giving me some content pillars, even some fun ideas that I might be able to execute on social.

If I go back into the notebook over here, one of my favorite things is actually an audio overview. Copilot will essentially put together a podcast that gives you a fun, interactive way to experiment with your content. Let's listen to a couple seconds.

COPILOT AGENT 1: Hey there. Welcome, everyone, to today's overview on urban mobility innovation. I'm really excited to dive into a fresh concept that's been turning heads.

COPILOT AGENT 2: Absolutely. It's a great day to explore new ideas.

CALLIE AUGUST: This is just simply amazing. And the best part is your notebook will automatically update over time as documents in the notebook are updated.

Back to you, Satya.

(Applause.)

SATYA NADELLA: Thank you very much, Callie. I love those audio overviews. I never realized that these podcast formats do really help you consume some of the more dense information, especially in the context of your project. I never thought of that application, but it turns out, every day, think about your audio overviews or SharePoint being used in a podcast format is just fantastic.

That gave you a little bit of a flavor for how you use something like Copilot for knowledge work, but I want to end by talking about one other domain that I think is

probably going to be the place where some of the most profound impact of AI will be felt, which is science.

If you take any business of ours, it doesn't matter, you can be in consumer packaged goods, or you can be in pharma. You can be in manufacturing. You can be in the business like us, in building data centers. Every one of these also has, essentially, some scientific loop that is driving innovation. And so, to that end, we have taken the inspiration from what we're doing with Copilot and said, what if we built essentially, an AI tool to accelerate scientific discovery?

And that was the motivation for Azure Discovery that we just launched in May at our developer conference. And what it does is it takes models that were built for physics or material science or chemistry, and then puts the same type of multiagent frameworks. In fact, one of the cool innovations that came out of this, that now is going to be part of even Foundry, is Graph RAG or graph-based retrieval, augmented generation, which I think is a good way for us to even think about scientific literature being available as grounding as you think about the new hypothesis creation for any scientific process.

That's Microsoft Discovery, which is a new tool chain and a new set of tools. And in fact, two examples of it that we ourselves used, one was when we built out a new – basically an electrolyte candidate, which has 70% less use of lithium for a battery. That's just, if you think about one of the challenges for us is the rare earth minerals and batteries, the ability to come up with complete new material science to help us there. So, we have that.

In fact, we went end to end. We worked with one of the national labs in the United States and fabricated it all the way. This is a material that didn't exist before, and we were able to use this discovery tool to build out and finally synthesize it completely.

Another one, which we recently did, was a new coolant for immersion cooling in a data center. Especially with some of the AI accelerators. AI cooling is not sufficient. You need to go to immersion cooling. And if you think about the immersion cooling, it has to be good with the right dielectric constant, such that you can have the right insulation for the electronics. And we were able to build out, effectively, again, a new coolant that we, ourselves, will use in our data centers. But these are two examples of how you can start thinking about using AI tools to create new materials, new coolants and so on.

Now, the one place where this is all going is quantum. Today, a lot of what we do with AI discovery is using HPC to essentially do the simulation, and AI to accelerate the simulation and learn the language of nature. Quantum is the ultimate breakthrough when it comes to the language of nature.

To that end, we are continuing to push forward with our quantum program. We have had major breakthroughs with this new topological quantum chip, Majorana 1. In fact, it has a deep connection to some of the innovation and contributions of our lab at the Tech University of Delft, right here in terms of some of the fabrication IP, as well as the

measurement protocols, which are critical, that got built right out of here. Let's roll the video to give you a flavor for Majorana 1.

(Majorana 1 video segment.)

SATYA NADELLA: That was a little bit of a tour of the full tech stack behind this open, agentic web. The great news here is that the maturity of the platform is there, the ambition of the developers is there. And now, it's a question of us being able to use this, ultimately to realize the mission that we collectively have, which is to empower every person and every organization right here in Netherlands to be able to achieve more.

Thank you all very, very much. Enjoy the rest of the show. Thank you.

(Applause.)

END